CHAPTER

7

DESIGNING A WAN INFRASTRUCTURE

After reading this chapter and completing the exercises, you will be able to:

- Design a Routing and Remote Access (RRAS) solution to connect various locations
- Design a demand-dial routing strategy
- Design an implementation strategy for dial-up remote access using RRAS
- Design a virtual private network (VPN) strategy
- Design an RRAS implementation solution for dial-up remote access

In this chapter, we explore the challenges of designing for wide area network (WAN) routing using Windows 2000 Routing and Remote Access (RRAS). First, we briefly look at the capabilities of RRAS in Windows 2000, and then we provide an overview of the virtual private network (VPN) technologies available in Windows 2000. We also determine what is needed to design a routing solution for connecting networks at multiple locations, utilizing demand-dial routing when appropriate, and VPNs for security. Shifting gears slightly, we then examine options for dial-up remote access for clients accessing a corporate network and consider various options to provide secure communications in our dial-up solutions.

WINDOWS 2000 RRAS BASICS

Windows 2000 RRAS is a software router and dial-up remote access server that provides a variety of routing services. These services include multi-protocol LAN-to-LAN, LAN-to-WAN, VPN, and **Network Address Translation (NAT)**. The latter is an Internet standard that enables a LAN to use one set of IP addresses for internal traffic and another set of addresses for access to an external network, usually the Internet.

The following facts about RRAS are important for use in network design:

- RRAS supports nonpersistent connections through the use of **demand-dial connections**, which are network connections initiated when data needs to be forwarded. Demand-dial connections are usually terminated when there is no traffic. This is ideal for a circuit-switched WAN link.
- RRAS reduces traffic and provides for the design of secure solutions for communicating over private and public networks using router authentication and encryption of data between routers.
- Computers running RRAS can be configured to perform **Internet Control Message Protocol (ICMP)** router discovery, by which a host can discover a router automatically, in spite of not having a default gateway configured in its TCP/IP properties.
- RRAS can isolate a private network, restricting the flow of incoming and outgoing traffic through the use of IP filters.
- RRAS supports multiple transport protocols, including TCP/IP, IPX/SPX, and AppleTalk.

RRAS supports the following routing protocols: Routing Information Protocol (RIP) for IP, RIP for Internetwork Packet Exchange (IPX), Open Shortest Path First (OSPF), Internet Group Management Protocol (IGMP), and Service Advertising Protocol (SAP) for IPX. In addition, when combined with other network services, RRAS can provide the Resource Reservation Protocol (RSVP) for use with Quality of Service (QoS) activities, a reduction of undesired traffic, and router authentication and encryption of data.

THE ROLE OF VPN PROTOCOLS IN ROUTING AND DIAL-UP SOLUTIONS

A VPN involves the connection of a network or a single client computer to another network over an intervening network, which can be the Internet. VPNs are implemented to provide security over unsecured networks. Windows 2000 supports two protocols for establishing a VPN tunnel: **Point-to-Point Tunneling Protocol (PPTP)** and **Layer 2 Tunneling Protocol (L2TP)**. Both protocols are based on **Point-to-Point Protocol (PPP)**, which is a dial-in connection protocol not directly associated with tunnels.

The following facts about these protocols are important as you create your routing solution:

- PPP is a standard method for encapsulation of point-to-point network traffic that defines packet boundaries, identifies the protocol of the encapsulated packet, and includes bit-level integrity services. **Serial Line Internet Protocol (SLIP)**, the predecessor protocol to PPP, had serious limits, particularly the lack of protocol identification, bit-level integrity services, and security. SLIP is not used in our VPNs. PPP supports authentication protocols.
- PPTP is an Internet-layer protocol that encapsulates PPP frames within IP datagrams using Microsoft Point-to-Point Encryption (MPPE). This encapsulation is done before it transmits the packets over an IP internetwork. It requires an IP-based transit internetwork and each encrypted frame can be an IP datagram, an IPX datagram, or a NetBEUI frame with a Generic Routing Encapsulation (GRE) header. It does not, on its own, support tunnel authentication; **IPSec**, the new set of protocols for IP security, can provide tunnel authentication for PPTP.
- L2TP is based on Cisco's Layer 2 Forwarding Protocol and PPTP. It is used to create an encrypted, authenticated tunnel that requires IPSec transport mode for encryption. This means that you will need to install machine certificates on the VPN client and server if your IPSec configuration requires certificates. View the Windows 2000 Server Help topic "Machine certificates for L2TP over IPSec VPN connections" for more information. It does not require an IP-based transit internetwork, only a packet-oriented, point-to-point connection. Therefore, L2TP can run over IP, Frame Relay permanent virtual circuits, X.25 virtual circuits, or ATM virtual circuits. Both the source and destination hosts must support these protocols.

Designing an RRAS Solution to Connect Locations

In this section, we first consider the conditions that indicate the need for a routing solution and what information will help you to create a successful design. We then look at designing a functional RRAS solution for connecting multiple sites. Finally, we work on enhancing the design—first for security and then for availability and performance.

Designing a Functional Routing Solution

By now in your design process, you have determined that your network needs one or more routers. You have also gathered information about the distribution of client computers and the services that they must access in the network as well as the routing protocols in use on the network. However, you have not yet determined some of the finer issues surrounding routers. We discuss them next.

Business and Technical Needs for Routing

When you performed your business and technical analysis for your network, you gathered a great deal of information. Of the information you gathered, what points out the need for a routing solution? If you have one or more of the following conditions, the need for routers is established:

- Your private network includes networks at multiple geographic locations that must interoperate.
- You need to connect multiple network segments of different physical network technologies.
- You have multiple network segments and a need to limit the traffic.
- The technical environment supports industry standard routing protocols, including RIP, OSPF, and/or IGMP. The latter is not a true routing protocol, but is listed with the routing protocols in RRAS. (There's more on IGMP later in this chapter.)
- The security needs require router authentication and data encryption.

Once you determine that a routing solution is needed, you must gather other information required for design of an appropriate routing solution for private network connectivity. The information required includes the following:

- The number of locations
- The number of hosts and their distribution among locations
- The routing protocols supported by the existing network or by the design
- Security requirements for the design

Router Placement

When deciding where to place routers in your network design, you will have two major placement choices: within the network and at the edge of the network. Your decision will be based on your need to localize traffic and/or maintain security.

Placing routers within the network involves determining what network traffic must be localized and where the "security" boundaries must lie. After determining these issues, you can draw a map of the network with router locations. When designing router placement within a network, you should do the following:

- Ensure that network traffic is isolated
- Create screened subnets to protect confidential data
- Place routers to enable communications between dissimilar network segments
- Place routers to enable communications between dissimilar transport protocols

Placement of routers at the edge of a private network should result from a need for the following:

- Remote locations to exchange network packets via a public network
- A private network that is isolated from the public network
- The exchanges of packets between dissimilar physical private and public network segments

After you identify these needs, you should draw a map defining the location of routers at the edge of the network with interfaces to the public and private networks. Figure 7-1 is an example of such a drawing.

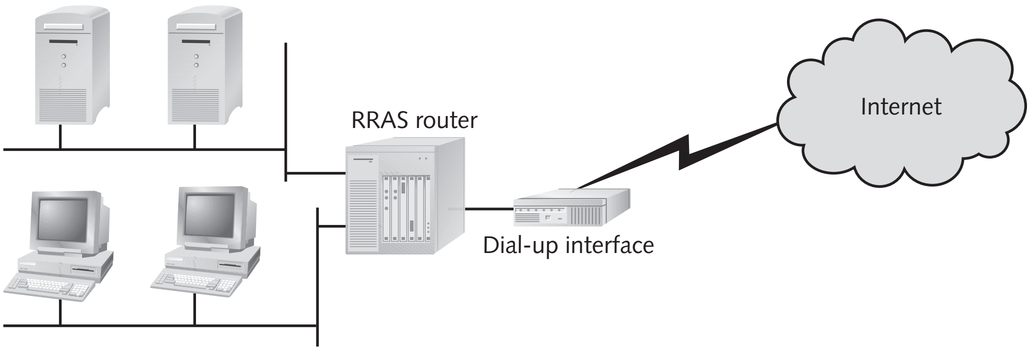

Figure 7-1 Router placement at the edge of a private network

Note

Although we do not always show the hardware devices needed to actually connect to the WAN, remember to take this into consideration in your design. Your router will be placed between the device and the private network. The device often acts as a bridge. Although they share commonly used names, the devices are unique to each connection type. For T-carriers, the connecting device is a Channel Service Unit/Data Service Unit (CSU/DSU). These are actually two devices, usually packaged as a single unit.

The CSU performs protective and diagnostic functions for a telecommunications line while the DSU connects a terminal to a digital line. (Imagine a very high-powered and expensive modem.) For an integrated services digital network (ISDN) connection, it is one of several specialized devices, generically referred to as an adapter; for a cable connection and a digital subscriber line (DSL) connection, the unique connecting devices are called modems.

Integrating Routers into the Existing Network

When integrating a router into the existing network, the designer must consider the configurations of the interfaces of the router. A router must have an interface on each

network it is connecting. Each interface is configured to work with the network to which it connects. Router configuration falls into three categories:

- Interface address and subnet mask
- Interface data rate and persistence
- Interface security

An interface address and subnet mask are required on each interface on an RRAS router. The address must be within the range of addresses of the network segment to which the interface is attached, and the subnet mask of the router interface address must match the subnet mask of that network segment.

Interface data rate and persistence are required configuration settings. For a router interface on the internal network, the data rate will be that of the attached LAN technology, such as 100 Mbps Ethernet, and it will be considered a persistent connection. An interface connecting to a public network may use a variety of technologies, including LAN and one of several demand-dial technologies.

A demand-dial technology is one that is initiated when data needs to be forwarded, and it can be terminated when there is no traffic. The data rate will be determined by the underlying technology. RRAS implements demand-dial connections through a logical demand-dial interface that represents the connection on the calling, or source, router. This software interface contains configuration information such as the address to use (phone number or IP address), port to use, methods of authentication and encryption, and authentication credentials.

Note

A demand-dial connection is also referred to as dial-on-demand (DOD). In fact, although used less frequently than the term "demand-dial," you will see this abbreviation in Microsoft documentation and technical articles.

Your design might include a **demand-dial interface** on top of technology that is persistent, such as using a VPN tunnel over DSL to have the added security of VPN tunnel authentication. A demand-dial interface is the software used by RRAS to control a demand-dial connection. You may want to avoid connection charges, as in the case of an ISDN connection. Adding a demand-dial interface on top of an ISDN connection will limit the connection charges to times when the connection is actually being used.

Interface security is important, regardless of whether a router is within the private network or on the edge connected to the Internet. The network designer has several guidelines to keep in mind when working toward interface security:

- Configure each router interface for the required security specific to that interface. (The security requirements for the network directly connected to an interface determine the level of security required for that interface.)
- Consider using encryption on private network segments hosting confidential data.

- If your organization is located in whole or in part outside the United States, verify that the data encryption level that you plan to use is legal.
- Authenticate all routers connected to public networks.
- Encrypt all data transmitted between routers that connect sites over a public network.

Configuration of Routers

There are two GUI tools that are especially important to the administrator of a Windows 2000 server for WAN connectivity. They are Network and Dial-up Connections, available from My Network Places, and the Routing and Remote Access Management console of Windows 2000, which is available from the Administrative Tools menu. Network and Dial-up Connections allows you to configure the drivers and protocols for each network interface. The Routing and Remote Access Management console allows you to configure and manage all the routing features of Windows 2000, such as router interfaces, VPNs, and filters. It is your central tool for establishing WAN, LAN, direct connect, and VPN connections for clients in Windows 2000.

When configuring a Windows 2000 server to be an RRAS router, you will take the following steps:

1. Configure each network adapter through Network and Dial-up Connections, ensuring that the proper driver is installed for the adapter and providing the IP address, subnet mask, and DNS server IP address.
2. Enable Routing in the Routing and Remote Access Management console by selecting Configure and Enable Routing and Remote Access and completing the wizard or choosing to manually configure once you have enabled RRAS.
3. Configure static routes (if appropriate) through RRAS.

Tip

Try Hands-on Projects 7-1 and 7-2 for more information on RRAS.

Selection of Protocols

Windows 2000 RRAS is an improvement over the previous version, especially in the administration tools. Its routing protocols, RIP and OSPF, are worth a closer look because you may need to use a Windows 2000 server as a router and RRAS is an objective for the Microsoft 70-221 certification exam. Microsoft also lists IGMP with the routing protocols; therefore, we will also examine where IGMP might fit into your design.

RIP

Before we proceed with the design issues of using the RIP protocol in Windows 2000 RRAS, let's take time out for a third-party view of RIP. RIP is a somewhat controversial protocol—router experts are extremely reluctant to use it. Why? First, it is very slow to converge. Second, if used on a demand-dial interface, it will bring up the link every 30 seconds to advertise its route, which could be costly and require special configuration of the ISDN link. These two issues are more significant than the bandwidth usage, which might only be about two seconds of every minute.

Keep the following facts in mind when considering using RIP in a design and when studying for the Microsoft 70-221 certification exam:

- If your environment includes networks with Novell IPX/SPX, and your other subnets need to interoperate with them, you may use the RIP for IPX capabilities of RRAS, coupled with the SAP for IPX protocol on the routers that must forward the IPX/SPX packets.
- Novell 5.*x* is the first version to support TCP/IP, in which case you can route the IP packets using RIP for IP or OSPF. Microsoft's client for Novell does not support TCP/IP, but the Novell client products support TCP/IP. Therefore, you could have Novell servers with either of these protocols.
- RRAS actually supports RIP Version 2, which supports Variable Length Subnet Masks (VLSMs), a feature you are likely to include in your IP addressing design for the network because it allows you to subnet a single network address into several smaller networks. Microsoft suggests that designers include RIP for IP in their designs for small networks to have routers automatically update routing table information.

Following are design considerations that should affect your network decisions concerning RIP:

- Because RRAS calculates the hop count of static route entries to be fixed at two, RIP protocol's 15-hop limit is reduced to 14 hops.
- RIP automatically updates the routing table.
- Consider RIP if the routing information changes frequently.
- Consider RIP if it is the protocol of existing routers, and they will or must be included in the new design.
- Consider using RIP to create auto-static route entries for demand-dial interfaces.
- Consider RIP version 2 in a network design in order to take advantage of its features. These include multicast router updates, support for **Classless Inter-Domain Routing (CIDR)**, VLSMs, and password authentication between routers. CIDR is a method of public IP address allocation that replaces the older system based on classes A, B, and C. CIDR was created to slow down

the rapid depletion of public IP addresses by allocating addresses with more flexible sizes of the ranges of addresses allocated.

At this point, we will discuss the protocols that are integral to your routing solution.

OSPF

OSPF is more efficient and has lower overhead than RIP and adds load balancing and class-of-service routing. Because of the need to extensively plan for OSPF and its more difficult setup, Microsoft recommends OSPF to network designers working with medium-size and large networks to support the automatic update of routing information for unicast packets. Recall that a unicast packet is one that has a single, globally unique destination host.

You should consider OSPF for your network design if:

- Routing information changes frequently.
- Other routers that will remain on the network use OSPF.
- Redundant paths do or will exist between two subnets.
- There are a large number of subnets (more than 50).

A design that includes OSPF benefits from a hierarchical perspective, in which the routers are grouped into three levels. At the top is the OSPF **Autonomous System (AS)**, in which all OSPF routers in the internetwork are included, with all OSPF routers on directly connected network segments. The middle layer of the hierarchy includes one or more OSPF areas, which include the routers connecting contiguous network segments.

Routers on the border of an OSPF area are known as **Area Border Routers (ABRs)** and connect their areas to a backbone area to which all OSPF areas within the AS connect. The bottom of the hierarchy is the OSPF network, which is a single network segment connected to one or more other OSPF networks through one or more OSPF routers.

With this hierarchy in mind, you may have already guessed that an OSPF design lends itself to a hierarchical IP addressing design. Once you establish the design for your OSPF hierarchy, you can subnet your IP network address into a hierarchy that maps to the design's AS/area/subnet/host levels.

In addition to this hierarchical perspective, consider the following steps when designing an OSPF AS:

- Create a high-bandwidth network segment for the backbone area of the AS.
- Minimize traffic by creating a **stub area**, an OSPF area that does not advertise individual external networks. A stub area uses a default route (network ID 0.0.0.0 and subnet mask of 0.0.0.0) for communication with external networks.

- Avoid **virtual links**, which must be created when an ABR is not directly connected to the backbone, but can be connected through a transit area to the backbone. This is generally accepted to be a poor design or something that only happens to resolve crises and provide work-arounds during network modifications. For more information on virtual links, see Windows 2000 Server Help.

An effective OSPF area design depends on the strength of the AS design, especially the hierarchical IP addressing. If this has been done carefully, the following strategies can be used:

- Use hierarchical IP addressing to assign to all areas TCP/IP network IDs that allow only a small number of routes.
- Use hierarchical IP addressing and route summarization to assign the single route that needs to be advertised to an area.
- If you place multiple ABRs in a single area, have them all summarize the same routes.
- Create your design so that all traffic between areas crosses the backbone area.

There are additional OSPF configuration considerations for your design, including:

- Control the designated router (DR) and backup designated router (BDR) by configuring the least busy routers with a higher priority than busier routers.
- Use a password for all routers in the same OSPF area.
- Microsoft recommends that you limit the number of network segments per area to fewer than 100, although this is not a hard and fast number. The actual maximum can only be determined through tests and depends on many factors, including stability, memory/horsepower of routers, and the use of summarization, to name just a few.

Note

An adjacency is a special relationship between neighboring routers for the purpose of synchronizing routing information. When two or more routers are on the same logical network, adjacencies will exist only between the DR or BDR and other OSPF routers.

IGMP

As more and more organizations include applications such as Microsoft NetMeeting or Windows Media Viewer in their application portfolio, network designers and managers must figure out how to modify the network infrastructure to support these applications. Although, strictly speaking, IGMP is not a true routing protocol, Microsoft lists it as such in the Routing and Remote Access console when you select the New Routing Protocol action for the General node under IP Routing.

Microsoft recommends that network designers include IGMP in a routing design to enable RRAS to send IGMP membership report packets from a single-router private

network to a multicast-capable portion of the Internet because of the limited IGMP capabilities of Windows 2000 RRAS. You should keep this limited implementation in mind. If you need more sophisticated support for multicast, such as is needed in a private network with multiple routers, you will need to find a third-party solution, such as those offered by Cisco or Nortel.

IGMP is typically implemented on switches, while special multicast routing protocols are implemented in the routers. These special protocols include Distance Vector Multicast Routing Protocol (DVMRP), Multicast Extensions to OSPF (MOSPF), Protocol-Independent Multicast Sparse Mode (PIM-SM), and Protocol-Independent Multicast Dense Mode (PIM-DM). None of these multicast routing protocols is included with Windows 2000; hence, it cannot communicate multicast information to other routers.

Note

To learn more about these multicast routing protocols, point your Web browser to *www.ipmulticast.com/community/whitepapers/introrouting.html*.

Windows 2000 RRAS and TCP/IP work together in supporting IP multicast traffic. TCP/IP for Windows 2000 supports IGMP in the following ways:

- It is an RFC 2236-compliant IGMPv2 host, which means it should work with third-party routers that support multicasting.
- It supports the mapping of IP multicast addresses to MAC addresses for Ethernet and FDDI network adapters per RFC 1112. When using Token Ring network adapters, IP multicast traffic is mapped to the Token Ring address 0x-C0-00-00-04-00-00.
- On RRAS servers, support for the forwarding of IP multicast traffic is based on the entries in the TCP/IP multicast forwarding table (viewable in the RRAS console) and a registry setting: EnableMulticastForwarding of the data type REG_DWORD in HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters. A value of 1 enables multicast forwarding, while 0 disables it. The default setting is 1 when the RRAS service is enabled and configured.

Let's look more closely at the capabilities of IGMP in the RRAS service, in which IGMP support is available in IGMP router mode or IGMP proxy mode. These modes are implemented per interface on an RRAS server, so different interfaces can be in different modes. IGMP router mode enables the interface to forward IP multicast traffic. To that end, it does the following:

- Switches the mode of the interface to promiscuous mode
- Listens for IGMP host messages and sends IGMP messages and query messages to manage local subnet multicast group membership
- Maintains the TCP/IP multicast forwarding table

IGMP proxy mode enables the interface to act as an IGMP-capable IP multicast proxy for hosts on the IGMP router mode interfaces. As such, an interface in IGMP proxy mode forwards IGMP host membership reports that have been received on IGMP router mode interfaces. It also does the following:

- Adds a multicast MAC address for each group address registered by proxy to the network adapter table of MAC addresses (Ethernet and FDDI). If the network adapter can support listening to all required multicast MAC addresses, it will be switched to promiscuous mode. The IGMP proxy mode interface will pass all non-local IP multicast traffic to TCP/IP for multicast forwarding.
- Updates the TCP/IP multicast forwarding table so that all non-local IP multicast traffic received on interfaces in IGMP router mode will be forwarded over the IGMP proxy mode interface.

Figure 7-2 illustrates how a Windows 2000 RRAS router with two IGMP interfaces—one in IGMP router mode and the other in IGMP proxy mode—can be used to connect a private network with a single router to an IP multicast-enabled internetwork, such as the Internet's Multicasting Backbone (MBONE). To learn more about MBONE, point your Web browser to *www.cs.columbia.edu/~hgs/internet/mbone-faq.html#topology*.

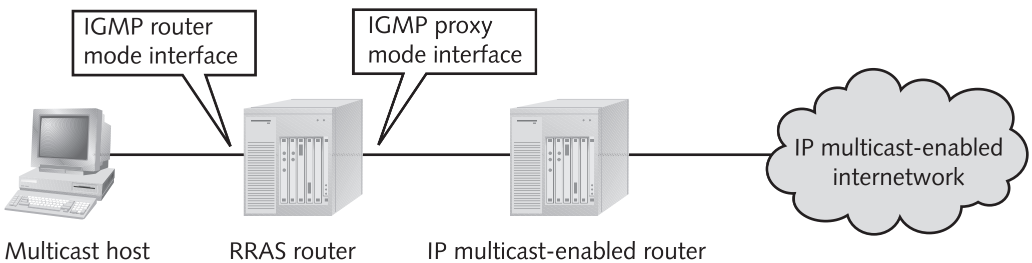

Figure 7-2 Connecting multicast hosts to an IP multicast-enabled internetwork

Integrating RRAS with Other Services

Windows 2000 RRAS is a service that depends on TCP/IP. It also works well with other Windows 2000 network services, allowing the following to happen:

- Integration with DHCP means that RRAS can use DHCP to allocate addresses for remote clients. Be sure to include DHCP Relay Agents in a routed design when DHCP traffic must be forwarded from DHCP clients to DHCP servers. As a best practice, Microsoft suggests placing DHCP Relay Agents on routers that only have connections to private network segments.
- RRAS will aid in forwarding DNS dynamic update traffic from RRAS clients, and RRAS also forwards dynamic updates of the WINS database from RRAS clients as part of normal packet forwarding.

- RRAS also integrates with Remote Authentication Dial-in User Service (RADIUS), allowing RADIUS to provide central authentication and record keeping for dial-in clients. (RADIUS is described in more detail later in this chapter.)
- RRAS can use IPSec for router authentication as well as the encryption of data transmitted between routers. Your RRAS design may need to use domain accounts and Kerberos Version 5 protocol certificates for router authentication.

Securing an RRAS Routing Design

To go beyond a design that provides simple routing functionality and to provide security, the designer must consider how to prevent unauthorized access to the network. This means that the design must not only protect the network from intrusion, but also protect the data being transmitted. RRAS supports security enhancements. We discuss each in turn next.

IP Packet Filtering

First, let's address TCP/IP filtering, a feature of Windows 2000 TCP/IP that's also available in NT 4.0 as TCP/IP security. This is not the filtering associated with routing. All NT 4.0 and Windows 2000 computers with TCP/IP have this feature, but it is turned off by default. Instead, TCP/IP filtering is used to allow an administrator to limit the incoming TCP/IP traffic for all IP interfaces. It's intended to be used to filter incoming traffic when another service or device (router, proxy server, or firewall) on the network is not filtering it. It can and should also be used in addition to filtering at the router; in which case, TCP/IP filtering at the individual computer level is your last attempt to thwart hackers.

TCP/IP filtering can be configured on any Windows 2000 computer through the TCP/IP properties of any network interface. Configuring this for one interface configures it for all interfaces on that computer. You can restrict network traffic based on TCP and UDP ports, as well as IP protocols. Figure 7-3 shows the TCP/IP filtering screen for a network interface.

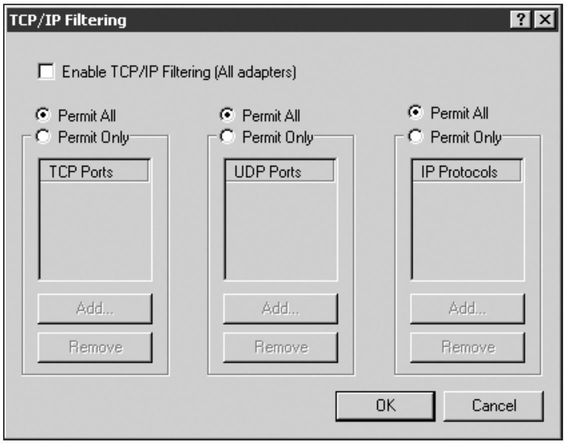

Figure 7-3 TCP/IP filtering settings

Now that you know what TCP/IP filtering is and where you may configure it, let's move on to IP packet filtering through the software router, RRAS. This is the type of filtering you are most likely to include in your network design because it is a more centralized approach to controlling traffic and is more configurable. It allows you to limit both incoming and outgoing traffic separately on each router interface. Figure 7-4 shows the Properties dialog box on the IP interface we have named "corporate" in our RRAS router. The Input Filters and Output Filters buttons give access to the dialogs for creating the filter list for this interface. Through RRAS IP filtering, you can control traffic based on the following:

- Source IP address
- Destination IP address
- IP protocol (TCP, UDP, ICMP) based on a protocol identifier, such as type and/or ports

These filters can be configured to *receive* all packets except those that meet the criteria specified in the filters or *drop* all packets except those that match the criteria specified in the filters. Test your design as thoroughly as time and money permit, because you do not want to lock down the routers unnecesssarily. You also do not want to leave security holes in your design. Testing will help you find these before the design is implemented. Try Hands-on Project 7-4 for more exposure to this issue.

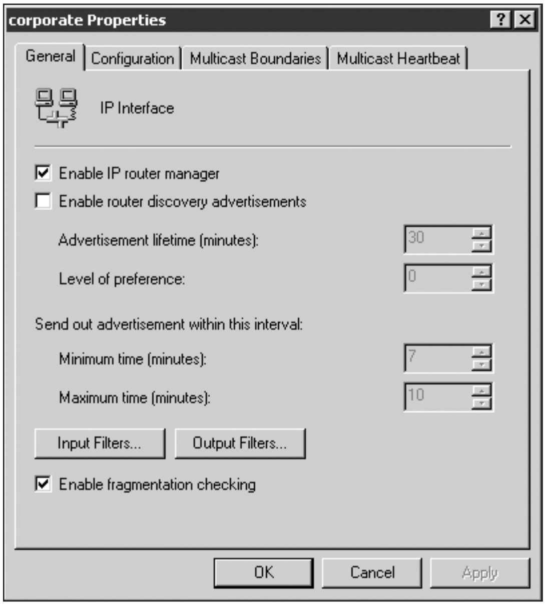

Figure 7-4 RRAS IP interface properties

Tip: If you are unsure of the ports in use on your network, use Network Monitor to capture each type of traffic, and then look for the port information in the headers of the packets.

IPSec

There are two major IPSec issues for the designer concerning routing. One is routing IPSec traffic that does not originate or end at the router; the other is using IPSec for router-to-router traffic.

When routing IPSec traffic that does not originate or end at the router, the router does not need to have IPSec settings that match those of the source and destination hosts. The router is simply forwarding encrypted packets. Getting IPSec traffic through a firewall, secure gateway, or proxy server is an entirely different problem, which is addressed in Chapter 9.

To enhance your RRAS design for security purposes, consider using IPSec to secure the router-to-router traffic. You can include IPSec on the routers in your design for router-to-router communication if you follow these rules:

- All routers involved must support IPSec.
- You must use machine-based certificates to authenticate the routers (more secure).
- You must use an Active Directory or public key infrastructure to issue the machine-based certificates.

Note

In order to use IPSec to authenticate the routers and encrypt data transmissions, Microsoft documentation states that the design must have only RRAS routers. However, many hardware routers do support IPSec and will support router authentication and encrypted data transmissions. This RRAS bias may show up on the test, but knowledge of both the RRAS software and hardware routers will help you before and after the test.

The last issue we need to address with IPSec is tunnel mode versus transport mode. Use IPSec tunnel mode to provide security in the form of authentication and encryption for router-to-router communications. When you use tunnel mode between routers, you must specify the IP addresses of the routers as the tunnel endpoints.

Note

IPSec transport mode establishes the level of security to be used between the router and any other computer. For this reason, you do not specify endpoints.

VPN Tunnels

Yet another option for authenticating routers and encrypting data in your RRAS design is to use VPN tunnels. In a design involving all RRAS routers, either IPSec or VPN tunnels can be used to authenticate the routers and encrypt data. VPN tunnels can be used between the routers in your design if you follow these rules:

- All participating routers must support VPN tunnels.
- Routers are authenticated with user accounts and/or the more secure machine-based certificates.

Your VPN protocol choices include PPTP and L2TP. PPTP will be your choice if you are including NT 4.0 RRAS routers or hardware routers that support PPTP. The encryption protocol used with PPTP is **Microsoft Point-to-Point Encryption (MPPE)**, which includes either 40-bit or 128-bit encryption.

Implementing 128-bit encryption in your design may require you to become familiar with government laws. If all or part of your company operates outside of the boundaries of the United States, you will want to check the laws of all countries involved, because prior to 2000, the United States had laws against exporting any products that contained encryption stronger than 56 bits. Although the laws were changed, the tides of international trade and intrigue flow in all directions; therefore, verify that your design does not break the laws of any country before you implement a high encryption protocol.

You can choose PPTP with MPPE for encryption if your design includes the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), MS-CHAP Version 2 (MS-CHAP v2), or Extensible Authentication Protocol Transport Level Security (EAP-TLS). MS-CHAP and MS-CHAP v2 are Microsoft's implementation of Challenge Handshake Authentication Protocol (CHAP), in which the remote access server sends a challenge to the remote access client consisting of a session ID and a challenge string. In Version 1,

the client must return the user name and a Message Digest 4 (MD4) hash of the challenge string, the session ID, and the password. Version 2 provides stronger security than either CHAP or MS-CHAP, providing mutual authentication so that the server is authenticated as well as the client, and support for asymmetric encryption keys. EAP is an extension to PPP that allows for new authentication mechanisms to be used for validating remote connections at the transport level. Where earlier methods depended on a certain authentication protocol, EAP allows for a negotiation to determine the authentication protocol to be used for a connection. EAP-TLS is an authentication method that can be used with EAP. EAP-TLS requires that the client and server use certificates to perform mutual authentication.

You can choose PPTP with MPPE for encryption if the security requirements for your design will allow for user-based authentication, rather than machine-based authentication, or if no machine-based certificate infrastructure is available.

L2TP is your choice if your design includes Windows 2000 or hardware routers that support L2TP, in which case, IPSec is used to encrypt the data. IPSec in Windows 2000 supports several encryption levels, including 40-bit DES, 56-bit DES, or Triple DES (3DES) encryption. (Remember that depending on the laws of the United States at the time you implement your design, the highest level of encryption may not be available for export.) Windows 2000 IPSec uses machine-based certificates for authentication, which provides a higher level of security than user-level authentication.

You should consider using L2TP tunnels with IPSec for data encryption if the security requirements for the design are too high to permit user-based authentication, and you need the higher security of machine-based authentication. You also should consider it if an Active Directory domain, or some other source of machine-based certificates, is part of the design.

Whichever method you choose to create and authenticate your tunnel, you will also be faced with the issues of configuring a tunnel for persistent connections versus demand-dial connections. You can assume the following if you are configuring a tunnel between routers communicating over the Internet:

- If your router-to-router VPN tunnels exist over persistent (full-time) connections, each router only needs a single demand-dial interface.
- If configuring for an on-demand connection, the destination router must be permanently connected to the Internet, while the source router connects to the Internet using a dial-up link. The destination router will need a single demand-dial interface; the source router will need two demand-dial interfaces—one for connecting to the ISP, the other for the VPN.
- The source router must have two static entries in its routing table: a static host route to the ISP so that it can dynamically connect to the Internet and a static route to the destination router. Figure 7-5 illustrates a router-to-router VPN tunnel over a dial-up connection.

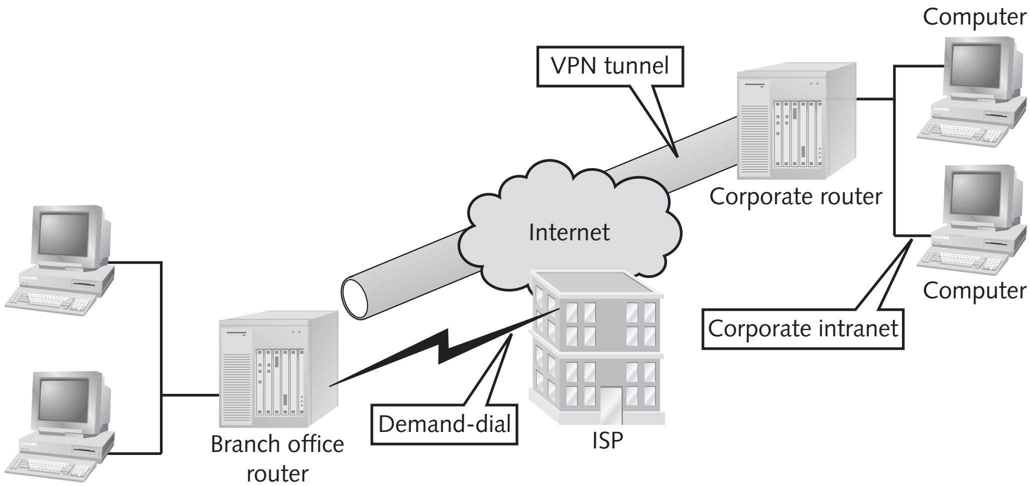

Figure 7-5 Router-to-router VPN tunnel

Note

For the nitty gritty details on configuring a VPN tunnel between routers over a dial-up connection, search Windows 2000 Server Help for the title "Router-to-router VPN design considerations."

Router Authentication

Yet another security enhancement for a routing design is router authentication, which guards data from being picked up by unauthorized routers. You can do the following for router authentication (note that the first two items in the list are likely to appear on your exam):

- RIP-for-IP passwords can be used if all routers use RIP and if you are capable of using "clear text password exchange" and "security" in the same sentence (yes, that is sarcasm).
- OSPF passwords can be used if all routers use OSPF and clear text password exchange is palatable to you. This is a big drawback, because if some of the routers in an area are third-party routers capable of more secure authentication methods, they will have to be "dumbed down" to use clear text passwords.
- You can use demand-dial authentication between routers that use demand-dial interfaces, in which case you have many choices, because demand-dial can use any RRAS-supported authentication protocol. Additionally, you can choose between one-way and two-way authentication. If you use one-way authentication, the calling router is authenticated with a predefined account and password, but it has a major drawback in that the destination router is not authenticated. Two-way authentication (a.k.a. mutual authentication) requires that both routers be authenticated using a predefined account and password. Two-way authentication does require MS-CHAP v2 on both routers.

- IPSec machine certificates provide perhaps the most secure router authentication method. Use them for router authentication in your design if all participating routers support IPSec and if you have high security requirements.

Screened Subnets

A screened subnet is used to protect a private network from the Internet, yet allow private traffic to be forwarded between intranet sites. Router placement is a significant issue when considering the use of screened subnets in your design, because the routers (or better yet, the firewalls) define the boundaries of the screened subnets. Basically, you place routers at the edge of the private intranet to create the screened subnets and you place routers between screened subnets to forward traffic.

Microsoft suggests creating screened subnets with the use of RRAS IP filters. However, they add the condition that you do this only if IP filters are adequate for the security requirements of the design *and* if the router is connected to the Internet and screens the private network from the Internet.

If your design includes the creation of screened subnets with firewalls (or proxy servers), you may still want to provide additional security by using RRAS routers with IP filters configured between the subnets of the private intranet, as shown in Figure 7-6. This also allows you to further restrict traffic to one or more of the screened subnets.

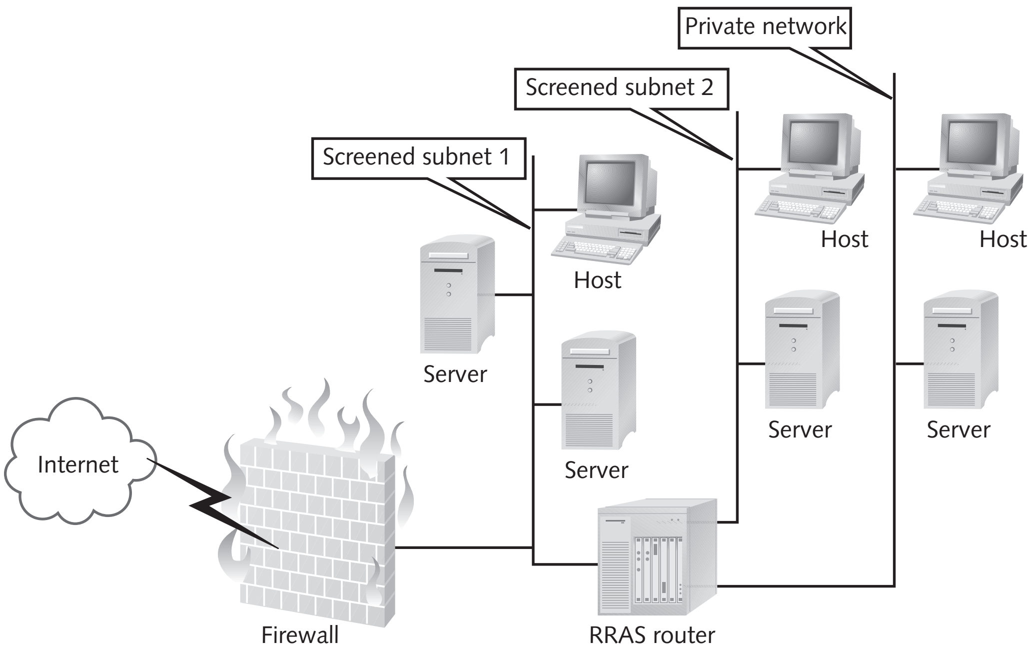

Figure 7-6 A router between screened subnets

Note

In some of its documentation, Microsoft uses "screened subnet" and "DMZ" interchangeably; however, in more common usage, what is pictured in Figure 7-6 is a screened subnet, and a DMZ would be a subnet isolated between two firewalls. For read more on this, point your Web browser to *www.microsoft.com/technet* and search for the article "Data Security and Data Availability in the Administrative Authority." Microsoft presents a view more in line with the industry in their Exchange Technical Notes article, "Best Practices for Developing ASP-Hosted Exchange 2000 Services."

Static Routes

As a designer, you will consider using static routes selectively in your design. The gains are in both security and performance. Although this section is about enhancing your routing design for security, to provide a more complete static route discussion, we will include both the security and performance aspects of static routes, as well as the use of static routes with demand-dial interfaces.

When you consider static versus dynamic routing, you might only be thinking of the work involved in manually adding routes to a routing table as opposed to allowing the routers to communicate with each other and build their own routing tables. In a small, stable network with few subnets, you may choose static routing to keep the design simple and to avoid the overhead of routers talking to each other. But then you might jump to the conclusion that static routing is simply too much work for your medium-size or large network and discount it out of hand. However, it is not the simple choice you may believe it to be.

First of all, each router should be considered from the perspective of its placement and function on the network. Then you should consider static routing for your design if one or more of the following apply:

- You wish to reduce the router-to-router traffic of dynamic routing protocols.
- Your security requirements demand that routing tables not be communicated on the network. In this way you can avoid being open to malicious acts, such as someone sending bogus routes to your routers, leading to denial of service (DoS) attacks.
- Manual updating does not take more time than administrators can afford to devote.
- The network routing table information is very stable.
- You require a demand-dial interface and need to add a default route to the interface.

The value of reducing router-to-router traffic can be significant on a segment that is already contending with a growing amount of traffic. Reducing traffic is always an admirable goal, as long as it does not require a prohibitive amount of administrative attention or hurt another design requirement, such as security.

Security requirements are a notable reason for using static routing. Avoiding dynamic routing protocols results in a higher degree of security because the routes are not communicated between routers and thus avoid the possible capture of the routing traffic, which would reveal information about your private network.

You should strive for manual updating time that is not excessive. This is an intangible that network designers and managers will have to determine among themselves, balancing the benefits of static routing, such as security and reduced traffic, with the cost in staff time for manual updating.

A stable network routing table will minimize the administrative cost of static routes. Add a default route to the demand-dial interface in order to have all IP packets with destinations outside the private network forwarded to the demand-dial interface. This static entry only needs to be added once, but consider that all traffic not intended for the internal network will be forwarded to the demand-dial interface. Try Hands-on Project 7-5 for more exposure to this issue.

A static route connecting two single-subnet networks would be fairly simple, but adding a static route to an RRAS router when the internal network has more subnets can be a little tricky. Microsoft TechNet article Q178993 gives the necessary steps that must be taken for such a scenario. Figure 7-7 shows such a scenario in which a branch office is connected to a central office network. Because of the need for RRAS Server 1 to route to all three subnets, the RRAS routers would have the static router configuration shown in Table 7-1.

Table 7-1 Static Routes for Figure 7-7

	RRAS Server 1	RRAS Server 2
Destination	10.30.0.1	0.0.0.0
Network Mask	255.255.0.0	0.0.0.0
Gateway	1.1.1.1	1.1.1.1
Metric	1	1

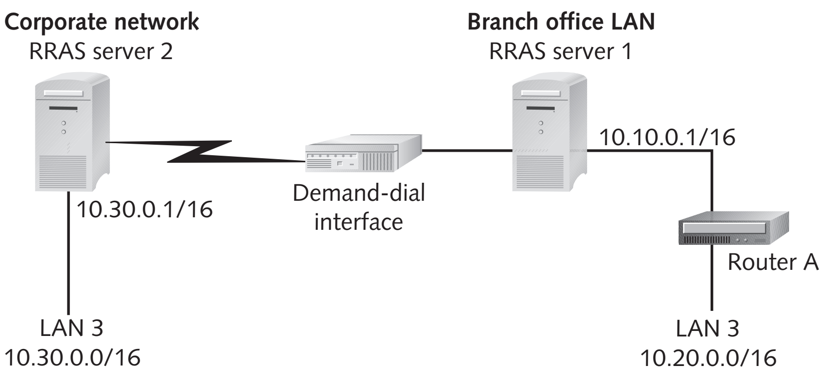

Figure 7-7 Using static routes in a branch office network design

Of course, you might also find scenarios in which static and dynamic routing can be combined. In that case, you can use auto-static route entries, which is a hybrid of static route entries and dynamic routing table entries. These static routes are added to the routing table automatically at scheduled intervals by the RIP protocol. This is a solution for adding static routes to a remote RRAS server connected to the corporate network with a demand-dial interface. Auto-static routes are only supported by RIP-for-IP, RIP-for-IPX, and SAP-for-IPX. OSPF does not support auto-static route entries.

To bring all this information together, let's consider a scenario in which a branch office is connecting to a corporate network through an RRAS server. The RRAS server at the branch office is configured to update its auto-static route entries once a day at 2:00 AM. At that time each day, the remote router initiates a demand-dial connection to the home office, deletes existing auto-static route entries matching the updates received, and then adds the new auto-static route entries.

Why use auto-static routes? With auto-static routes, unreachable networks will not cause the router to initiate the demand-dial connection. However, your design must include procedures for keeping the static route information up-to-date when subnets are added to the network on the other end of the WAN connection.

Improving Routing Availability and Performance

A routing design can be enhanced for availability and performance with four simple strategies: dedicated routers, persistent WAN connections, redundant WAN connections, and redundant routers. These strategies work independently and have wonderful synergy when combined. Let's explore them.

Dedicated Routers

The best strategy for improving routing availability and performance is to use dedicated, single-purpose routers (hardware routers). However, if you are using a Windows 2000 computer as your router, you will enhance availability and performance by limiting its purpose to routing. As our favorite "router guy" says, "I want my servers to serve and my routers to route." This enhances availability by eliminating the mishaps that can occur when unstable applications and even OS components (like the print spooler) go haywire (an important technical term) and cause or necessitate a reboot. This strategy optimizes performance by eliminating the competition for resources that a multi-purpose server presents.

Persistent WAN Connections

Another important enhancement for router availability and performance is to select WAN connections that are persistent. A persistent WAN connection is one that is always active. Availability is enhanced because you avoid having to initiate the connection when the router receives data with a destination across that particular connection. Performance is enhanced because there is no delay while a connection is established.

Redundant WAN Connections

Providing redundant WAN connections enhances availability by avoiding downtime if one of the connections fails. This strategy also enhances performance if the redundant connection is not simply held in reserve as a "failover," but used to distribute the traffic load across the redundant WAN connections.

Redundant Routers

Redundant routers are insurance against downtime due to the failure of a single router. The use of redundant routers also can enhance performance by distributing the traffic load across the redundant routers. If your network design requires a high level of availability and performance, and your company has the money to fulfill this requirement, your design could combine all four strategies and have dedicated, redundant routers and persistent, redundant WAN connections.

Designing an RRAS Solution for Dial-up Remote Access

Windows 2000 RRAS provides the dial-up environment for remote users connecting to private networks. Like Windows NT, dial-up clients can be limited to accessing resources on the remote access server itself, or they can be allowed to access other resources on the organization's internal network. In this section, we examine the technologies and tools available in such a design.

In the process, we look at VPN technologies for use with remote access clients. Following the VPN section, we explore the use of something new to Windows 2000 RRAS—remote access policy, which is critical to a functioning remote access design. Along the way, we look at how each of these features can contribute to a functional design, and later we explore how to combine RRAS and other Windows 2000 technologies to make a design more secure. Of course, security often comes at the price of performance, so we end our discussion with ways to enhance a remote access design for performance and availability.

Designing a VPN Strategy for Remote Access

Several years ago, dial-up service for remote users was commonly hosted within an organization, with the connections made over costly phone lines. At that time CIOs (if they existed), network managers, and operations managers would never have allowed the use of the Internet for an organization's site-to-site network communications. Now, VPN technologies provide cost-effective, secure solutions for connecting remote locations to a private intranet using a public network, usually the Internet, as the backbone.

In spite of the cost savings, decision makers still question the reliability of the Internet for site-to-site WAN connections as well as the quality of transmissions, particularly the characteristic Internet problem of jitter. Jitter does not hurt Web page browsing or

e-mail communications, but it can render time-dependent communications such as streaming video useless.

A voluntary, as opposed to compulsory, VPN tunnel occurs when a client establishes a VPN over an Internet connection, with a server on or at the edge of the private network serving as the endpoint. In this case, the client must support the tunneling protocols, and no intermediate server can be used to create and maintain the tunnel with the corporate server. We call this tunnel voluntary because one of the endpoints of the tunnel is at the client computer (see Figure 7-8).

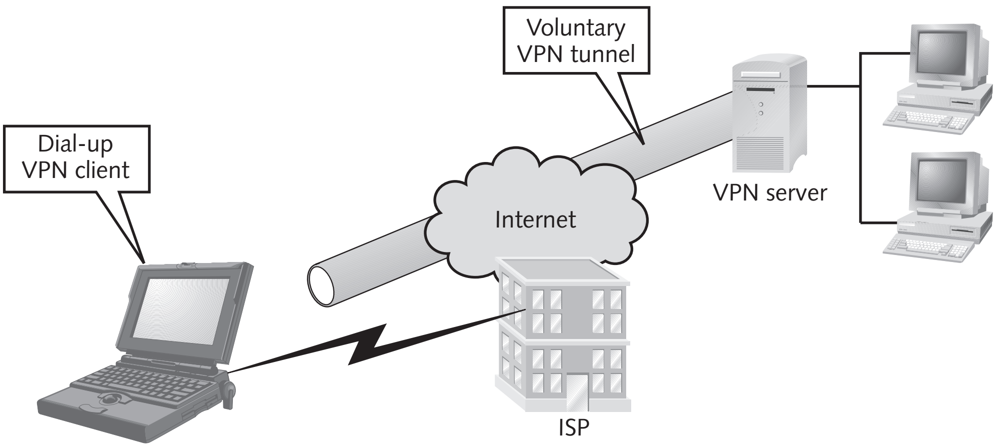

Figure 7-8 Dial-up VPN client using a voluntary tunnel

Another VPN connection type is the compulsory VPN tunnel, in which the client is not required to support the VPN protocols. The remote access server initiates the tunnel connections and supports the tunnel protocols. User authentication is required, but the client is not involved in creating the tunnel. The remote access server may also use RADIUS, which we will examine more closely in the next section. This has led to scenarios in which the remote access dial-up point is hosted at an ISP, in which case, clients first connect to the ISP and then connect via a VPN from the ISP to the remote network.

Figure 7-9 illustrates a scenario in which the dial-up client connects to the ISP where a compulsory tunnel carries traffic between the client and the VPN server in the private network.

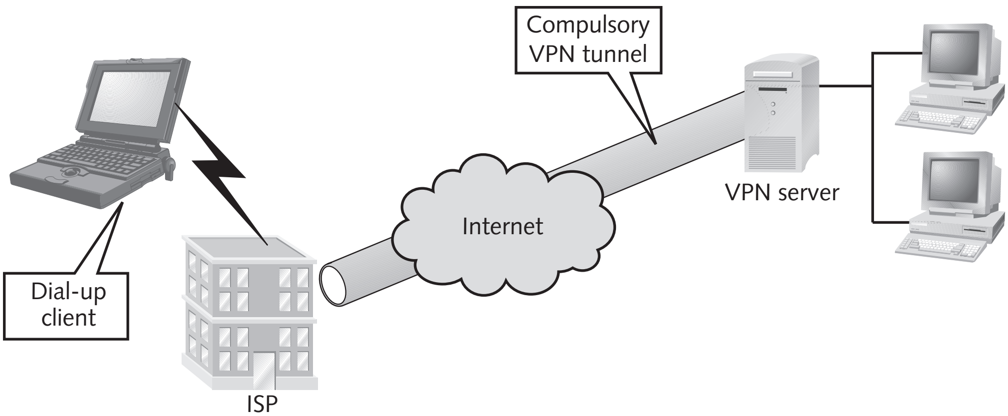

Figure 7-9 Dial-up VPN client using a compulsory tunnel

When to Use VPN for Dial-Up Access

When determining remote access design needs, you first must evaluate if a simple dial-up connection will suffice or if the security requirements call for a VPN. If the client is dialing directly into the private network, dial-up access without a VPN will suffice. A VPN should be part of your network design if the connection is being made over the Internet, the connections to the Internet (both client-side and server-side) can support the maximum expected load from client access, and the performance of the connections is acceptable.

The designer must also determine configuration settings needed for the VPN server, including:

- The tunneling protocol that best fits the design requirements
- PPTP ports and L2TP ports required for client connections
- Which user accounts are to be granted access
- What remote access policy settings are required

Note

The tunneling protocols available in this scenario are PPTP and L2TP, which were discussed earlier in the chapter.

Whichever VPN protocol you use, the routers between the endpoints do not need to support the tunnel protocols—they simply forward the packets. In addition, if clients can support VPN protocols, and you desire a high level of security, use voluntary end-to-end VPN tunnels. If the clients cannot support VPN protocols, use compulsory VPN tunnels. In addition, remember that you must determine the number of PPTP ports and/or

L2TP ports required for client connections to each RRAS server. The default number of 128 of each port type can be modified to fit the projected number of concurrent VPN connections.

As in Windows NT, Windows 2000 RRAS Server requires that dial-in users be granted access. You must determine which user accounts need to be granted access during the business and technical analysis, and then you determine how to grant access. We will move into this topic shortly. For now, understand that information gathered in the business and technical analysis will help you to determine the remote access policy settings to control many aspects of a dial-up session, such as user groups, tunnel types, and client IP.

Strategies for Using Remote Access Policies

We are sure that the typical student of this course has either spent years working with networks and/or long months learning about networking, operating systems, and Microsoft products before approaching the subject of network design. We still take the risk of including review information from time to time that could, strictly speaking, be considered part of the prerequisites for this course. We do it because, sometimes, a topic is simply too important for the success of the student both as a network designer and as a test taker. **Remote access policies**, a set of conditions and connection settings used to grant remote access, fall into this category.

If you are working to earn your Windows 2000 MCSE and have taken the four core exams, you also have encountered this topic already. However, beginning in Windows 2000, it is entirely new and not easy to master; therefore, we will take the plunge and present you with some essential information to help you use remote access policies in your design. This could be career-saving knowledge, because it is easy to misunderstand remote access policies and create a design that is unusable, allowing no remote clients to connect, or that is too open to meet the security needs of the design.

Here are a few basic facts and rules for working with remote access policies:

- Remote access policies are stored on the remote access server, not in Active Directory. It does not depend on Active Directory, although there can be a relationship, which we will discuss a little later. Having remote access policies reside on a Windows 2000 remote access server means that policies can be varied according to the capabilities of the server and the communications links it is using.
- No remote access policy means no remote access. RRAS includes a default remote access policy. If you delete this without creating a new one, your remote access server is going to get very lonely.
- Dial-in permissions set in the user account properties will override remote access permissions.

Sounds simple, huh? Actually, remote access policies are made up of simple parts, but a lot of them. We discuss them next.

Remote Access Policy Components

The many possible settings of remote access policies are organized into three components: conditions, permissions, and profile. We take a brief look at each of these components and the settings they can contain before we look at how all these pieces work together to permit or deny remote access.

First, we look at conditions. Remote access policy conditions are a list of settings, including phone numbers, client IP address, day and time restrictions, protocol, tunnel type, and user groups. The complete list of conditions can be viewed in the Properties of a Remote Access Policy in the Routing and Remote Access console.

Now, we consider permissions. The remote access policy permissions settings are evaluated together with the dial-in permissions on a user's account in Active Directory. The remote access policy user permissions setting is very simple, with just two choices: grant remote access or deny remote access. However, its relationship with the permissions set on the user account settings and groups defined in the conditions makes it more complicated.

The user account dial-in permissions will override the remote access policy permissions. If the user account dial-in permission is set to allow access, the user will be permitted, even if permission on the remote access policy is set to deny access. The reverse is also true: If the user account dial-in permission is set to deny access, the user will not be able to dial in, even if the remote access policy permission is set to grant access.

Note that the "Control Access through Remote Access Policy" option is available *only* in a native-mode domain; it is "grayed out" when a domain is in mixed mode. If this setting is selected, the dial-in permission depends on the remote access policy permissions. In a mixed-mode domain, only the allow access, deny access, and callback options are available as user account dial-in settings.

And finally, we take a look at profiles. Each remote access policy has a set of profile settings that are grouped into several categories. These include Authentication, Encryption, Dial-in Constraints, IP, Multilink, and Advanced. Authentication allows an administrator to determine what (if any) authentication method is used. Encryption settings control the encryption levels that will be accepted. An administrator may select any or all of the following settings: No Encryption, Basic, Strong, and Strongest.

If No Encryption is the only encryption level selected, clients using data encryption will not be allowed to connect. The Basic setting allows IPSec 56-bit DES for L2TP over IPSec-based VPN connections or MPPE 40-bit data encryption for dial-up and PPTP-based VPN connections. Strong means IPSec 56-bit DES for L2TP over IPSec-based VPN connections or MPPE 56-bit data encryption for dial-up and PPTP-based VPN connections. Strongest uses MPPE with a 128-bit key for dial-up and PPTP-based VPN. For L2TP over IPSec-based VPN connections, Strongest uses triple DES (3DES) encryption. The Strongest option is only available on North American versions of Windows 2000.

Note

Dial-in Constraints has a large number of settings. These include options to disconnect if idle, restrict the maximum session time, restrict access to the days of the week and time, restrict dial-in to a specific number, and a raft of settings to restrict dial-in media (ADSL-DMT, Ethernet, IDSL-ISDN, and so on).

The IP settings of the profile allow an administrator to control the IP address assignment policy as well as apply IP packet filters to the connection. Address assignment policies include "Server must supply an IP address," "Client may request an IP address," and "Server settings define policy." The IP packet filters can be configured separately for traffic in each direction.

Tip

When you do Hands-on Project 7-6, take the time to examine the many conditions and profile settings.

Now that you have seen a summary of the remote access policy settings, let's look at how they work together, which is what occurs when a dial-up connection or VPN is initiated to an RRAS server.

Remote Access Policy Evaluation

When a remote access connection is attempted, the remote access policy and the user dial-in permission are evaluated to determine if the connection will be permitted. The evaluation follows these steps:

1. If there is no remote access policy for the RRAS server, no connection will be permitted. RRAS compares the conditions of the remote access policy or policies to the conditions of the connection.
2. If no policy has a condition that matches, the connection is denied.
3. If the conditions of a policy match, no other policy for the RRAS server is evaluated; the policy with the matching condition is the only one used, and evaluation continues with that policy.
4. RRAS checks the user account's dial-in permissions.
5. If the user account permission is set to deny access, the user is denied access.
6. If the user account permission is set to allow access, the user is granted access and the profile for the policy is evaluated and applied.
7. RRAS matches the connection setting to the settings of the user account and the policy profile.
8. If the settings match, the connection will continue.
9. If the settings fail to match at any time, RRAS disconnects the client.

Creating Remote Access Policies

To create remote access policies, you will need to know the security requirements and other remote access requirements and restrictions, such as the type of service available for the connections. With the information in hand, the administrator will do the following:

1. Configure the user account dial-in settings to either allow access or control access through remote access policy.
2. Create a policy in RRAS.
3. If control access through remote access policy was selected in the user account properties, you then grant permissions in the remote access policy and add groups under conditions, if you wish to control access through groups.
4. You then modify the conditions of the policy to match other requirements, such as time of day, user groups, and tunnel type.
5. You then modify the profile of the policy if there are design requirements that match the many profile settings, such as connection time limit, authentication requirements, encryption requirements, and type of media.

Security-Enhanced Dial-Up Designs

To create a secure remote access design, you must look at all the available features of RRAS and protocols and services that can be integrated with it to provide a more secure design. We will begin with security enhancement strategies available through the use of all the features, protocols, and services described in this chapter. We then consider strategies that include using RADIUS servers.

Selecting Protocols for Authentication and Encryption

Earlier, we discussed the protocols available for remote access authentication. Table 7-2 gives RRAS authentication protocols and the conditions under which you should consider each. With the exception of **Password Authentication Protocol (PAP)**—the predecessor to the others—these protocols provide encryption of the authentication exchange, but they vary in effectiveness.

Table 7-2 Remote Access Authentication Protocols

Authentication Protocol	Client Support/Benefit/Disadvantage
MS-CHAPv1	Supported by Windows 9*x*, Windows NT, and Windows 2000; encrypted authentication, but must be tweaked to be secure
MS-CHAPv2	Supported by NT 4.0 SP6a and later and Windows 2000 (documented as available with SP4, but not truly functional until SP6a); stronger authentication encryption than MS-CHAPv1
EAP-TLS	Gives smart card support to Windows 2000; adds per client hardware cost; needs RRAS and Active Directory

Table 7-2 Remote Access Authentication Protocols (continued)

Authentication Protocol	Client Support/Benefit/Disadvantage
CHAP	Because the challenge is passed, weak passwords can be broken; use with UNIX clients requiring encrypted authentication
SPAP	Encrypted authentication for remote access client using Shiva LAN Rover software, which is not very secure; use if Shiva server already is in place
PAP	Unencrypted authentication; use with clients that do not support any other protocol

Your design security requirements will lead you to select the best authentication protocol for remote access. And, as you can see, most authentication protocols encrypt the authentication traffic. You also want to consider the need to encrypt the data transferred after a user is authenticated for remote access. Both are selected through the remote access policy, and your selection of an authentication protocol affects your selection of data encryption protocols. The data encryption methods are divided between MPPE and L2TP over IPSec.

MPPE should be selected for remote access data encryption if the design includes MS-CHAP, MS-CHAPv2, or EAP-TLS authentication protocols. It also should be used if the security requirements call for user authentication, and there is no machine certificate infrastructure (like Active Directory).

L2TP over IPSec offers the most secure data encryption option for dialing into a Windows 2000 RRAS server. It requires a public certificate infrastructure; therefore, use this if your design calls for and can support L2TP tunneling and if a public certificate infrastructure exists, as is available with Active Directory.

Security Enhancement with Remote Access Policies

Now let's take what we know about remote access policies and boil it down to some simple strategies for making a remote access design more secure. Remote access policies are unique per server, so you will want to consider a separate set of policies per server. The exception to this is if the remote access servers are all using the same RADIUS server, in which case they all use the same set of remote access policies. Let's save the RADIUS discussion for a few minutes, just to keep this simple.

When creating remote access policies, consider the following strategies for remote access permissions:

- Allow or permit access by user: If you do this, you will modify the user account dial-in properties to allow or deny access. This strategy does not scale well, and should only be used when there are very few dial-in users.
- Allow or restrict dial-in access through a remote access policy in an Active Directory native-mode domain: In this case, the user's dial-in properties are

set to Control Access through Remote Access Policy, and the administrator can set remote access permissions to "on" by user or group.

- Allow or restrict dial-in access by a policy in a Windows 2000 mixed-mode domain: This is implemented by explicitly allowing or denying such in the user's dial-in properties. If allowed by user, the remote access policy can still deny by group through the remote access policy conditions.

You can select other remote access policy conditions (in addition to groups) to enhance security. You can set conditions that restrict access based on called-station ID, calling-station ID, day of the week and time, type of service the remote client requests, and type of tunnel.

Beyond the conditions, remote access policies also have profile settings that can enhance security. While the condition setting must be matched to accept the connection, the profile settings affect both the initial connection and parameters that will cause the connection to be terminated when they are exceeded. (Take a breath.) Some, but not all, of these settings are tied to security. Once again, day of the week and time appear under profile settings, but also appear as restrictions on the dial-in number, idle time, session length, and the actual dial-in media (such as modem, X.25, Ethernet, and many more). Profile settings also include authentication and encryption requirements, an IP address assignment method for the connection, and IP packet filters that will be applied to the connections.

Keep holding your breath—we're not quite done. You can configure multi-link settings and apply Bandwidth Allocation Protocol (BAP) settings to the connection. Then, if you need profile requirements that are not defined on the other tabs, you may use the Advanced tab of the Edit Dial-in Profile dialog box to specify additional connection attributes to be returned to the remote access server during a connection. The Add button on this tab sheet reveals yet another long list, such as several attributes that would apply to tunnels.

Enhancing Dial-Up Security and Administration with RADIUS

Remote Access Dial-in User Service (RADIUS) is an industry standard that has been around for a while. RADIUS centrally and securely authenticates remote access users who are outside the boundaries of a private network. It also maintains accounting logs of remote access usage that can be used for various tracking purposes, including charging for remote access service.

RADIUS fundamentals include some different meanings for familiar terms:

- **Authentication server**: This is the server that hosts the accounts database. In a RADIUS design in a Microsoft network, this would be a Windows NT or Windows 2000 domain controller. This server should be located on the same LAN as the RADIUS server.
- **Realm**: This is the entity containing the information for authentication (more global than authentication server). In a RADIUS design in a Microsoft network, the realm would be the NT4 or Active Directory domain.

- **RADIUS server**: This is the server that accepts authentication requests from a RADIUS client and authenticates the user accounts with an authenticating server. In a Windows 2000 network, this can be a third-party product (such as that from Shiva Corporation) or Microsoft's Internet Authentication Server (IAS). The RADIUS server should be located on the same LAN as the authenticating server.
- **Shared secret**: This is a password set by the administrator on both the RADIUS client and RADIUS server. This shared secret is used for communications between the RADIUS client and the RADIUS server, at which time each uses an algorithm that produces a hash of the password. The client passes this hash to the server, and the server compares the hash received with the hash that it produced locally. The RADIUS client also uses the shared secret to encrypt a remote access client password.
- **Remote access client**: This is not specifically a RADIUS term, nor a term that is modified; it is the remote access client requests for authentication at a remote access server (RRAS server) that are passed on by that remote access server as a RADIUS client.
- **IAS log file**: This is the file that holds the accounting information on a Windows 2000 IAS server. This file can be imported into a spreadsheet or database to access and query the information. Microsoft does not have a separate reporting tool for this information.

A functional RADIUS design must include a minimum of one RADIUS client and one RADIUS server. It also must include the following:

- A RADIUS client connection to the RADIUS server via a dial-up client connection, VPN client connection, a combination of these two, or a LAN connection
- A supported client remote access connection protocol, which can include TCP/IP, IPX/SPX, or AppleTalk
- Matching connection data rate, persistence, and security level between the RADIUS client and RADIUS server
- A default domain for the RADIUS server

Let's look at scenarios in which RADIUS is a solution for your network design. In the first scenario, consider an organization that uses an ISP for their Internet connections and that can host the dial-in connections for their mobile users. In this case, the design should include VPN tunnels between the dial-in clients and the ISP where a RADIUS client and authenticating server is located. The RADIUS client provides secure authentication with the RADIUS server in the private intranet.

Consider a second scenario. One company needs to give secure remote access to several partner organizations. In this case, a RADIUS client would be placed in each partner network, providing secure authentication of the remote access clients over a public network to the RADIUS server on the first company's private network.

Consider still another scenario. Various locations of a private intranet are connected through the Internet. In this scenario, the RADIUS clients are located in each of the regional and branch offices, and the RADIUS servers are located in the central office.

To enhance a RADIUS design for security, you must use remote access policies to restrict access, the highest level of authentication and encryption protocols available, RADIUS-shared secrets, IPSec machine certificates, and VPN tunnels. In addition, you must place RADIUS clients and servers within screened subnets.

You can enhance a RADIUS design for availability by having more than one RADIUS client and server in your design. Try Hands-on Project 7-6 for more expertise in this area.

Enhancing a Remote Access Design for Availability and Performance

It is probably a stretch to say that you can achieve availability and performance enhancements through remote access policies, but there are a few profile settings that can help. On the Dial-in Constraints tab of a remote access policy profile, the "Disconnect if idle for" setting can aid availability by limiting wasteful idle connections. The "Restrict maximum session to" setting also enhances availability, but may not be popular with users who have valid reasons for extended sessions. Although this may seem counterintuitive, restricting access to certain day and time—when carefully combined with user group settings—may also help availability.

We are striving to ensure that access is available to those who need the access to accomplish their work. If you can determine that certain groups should not be given access on certain days or at certain times, you are improving the availability for other groups who must have access at that time. You may also couple these settings with restrictions on dial-in media for performance.

The true availability and performance enhancements are available through hardware configurations, such as redundant WAN connections and persistent WAN connections, as illustrated earlier in this chapter. Using dedicated RRAS servers and redundant RRAS servers can also enhance availability and performance of a dial-up design.

Chapter Summary

- Windows 2000 RRAS is a software router and dial-up remote access server that provides a variety of routing services. These services include multi-protocol LAN-to-LAN, LAN-to-WAN, VPN, and Network Address Translation (NAT). The latter is an Internet standard that enables a LAN to use one set of IP addresses for internal traffic and another set of addresses for access to an external network, usually the Internet.
- A VPN involves the connection of a network or a single client computer to another network over an intervening network, which can be the Internet. VPNs are implemented to provide security over unsecured networks. Windows 2000 supports

two protocols for establishing a VPN tunnel: Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP). Both protocols are based on Point-to-Point Protocol (PPP).

- Strong designs consider the conditions that indicate the need for a routing solution and what information you will need to create a successful design. Related to this are the intricacies of designing a functional RRAS solution for connecting multiple sites. In all cases, you should enhance the design for security and then for availability and performance.
- Windows 2000 RRAS provides the dial-up environment for remote users connecting to private networks. Like Windows NT, dial-up clients can be limited to accessing resources on the remote access server itself, or they can be allowed to access other resources on the organization's internal network.

Key Terms

Area Border Routers (ABR) — OSPF routers that connect their areas to a backbone area to which all OSPF areas connect.

authentication server — A server hosting the accounts database for a RADIUS design.

Autonomous System (AS) — A group of routers on directly connected network segments that exchange routing information by using a common Interior Gateway Protocol, such as a system in which all OSPF routers in the internetwork are included, with all OSPF routers on directly connected network segments.

Channel Service Unit/Data Service Unit (CSU/DSU) — The hardware device used to connect a network to a T-1 or T-3 line.

Classless Inter-Domain Routing (CIDR) — A method of public IP addressing allocation that replaces the older system based on classes A, B, and C. CIDR was created to slow down the rapid depletion of public IP addresses, by allocating addresses with more flexible sizes of the ranges of addresses allocated.

demand-dial connection — A physical connection, such as a circuit-switch WAN link, that is initiated when a router receives packets to be forwarded to a destination across the WAN link.

demand-dial interface — The software component that recognizes the demand-dial connection on behalf of RRAS.

dial-on-demand (DOD) — An alternate term sometimes used instead of "demand-dial" in Microsoft documentation.

IAS log file — The file that holds the accounting information on a Windows 2000 IAS server.

Internet Control Message Protocol (ICMP) — A protocol by which a host can discover a router automatically, in spite of not having a default gateway configured in its TCP/IP properties.

IPSec — The new set of protocols for IP security built into IPv6 and implemented in the Microsoft IP in Windows 2000 and in later NT 4.0 service packs.

Layer 2 Tunneling Protocol (L2TP) — A protocol based on Cisco's Layer 2 Forwarding protocol and PPTP. It is used to create an encrypted, authenticated tunnel and requires IPSec for encryption.

Microsoft Point-to-Point Encryption (MPPE) — The encryption protocol used with PPTP that includes either 40-bit or 128-bit encryption.

Network Address Translation (NAT) — An Internet standard that enables a LAN to use one set of IP addresses for internal traffic and translates the internal addresses to a second set of addresses for access to an external traffic network (usually the Internet).

Password Authentication Protocol (PAP) — An Internet standard plain text authentication scheme included in Windows 2000 to allow clients to connect to non-Windows 2000 remote access servers and to allow non-Windows clients to connect to Windows RAS servers.

Point-to-Point Protocol (PPP) — A standard method for encapsulation of point-to-point network traffic that defines packet boundaries, identifies the protocol of the encapsulated packet, and includes bit-level integrity services.

Point-to-Point Tunneling Protocol (PPTP) — An Internet-layer protocol that encapsulates PPP frames within IP datagrams to be transmitted over an IP internetwork.

RADIUS server — The server that accepts authentication requests from a RADIUS client and authenticates the user accounts with an authenticating server.

realm — The entity containing the information for authentication (more global than authentication server).

remote access client — Dials in to a remote access server.

remote access policies — A set of conditions and connection settings used to grant remote access. Remote access policies are made up of many simple parts grouped into three components: conditions, permissions, and profile.

Remote Authentication Dial-in User Service (RADIUS) — An industry standard that offers centralized authentication of ISP or private remote access users. It is a security enhancement that also provides centralized accounting of dial-in connections.

Serial Line Internet Protocol (SLIP) — The predecessor protocol to PPP for sending IP packets over a serial connection.

shared secret — A text string that serves as a password between the RADIUS server and the RADIUS clients connected to it.

stub area — An OSPF area that does not advertise individual external networks. It is a portion of a network with a single entry and exit point that does not maintain routes to external Autonomous Systems.

virtual links — If a router designated as an ABR does not have a direct physical connection to the backbone, a virtual link can be created through an area that is connected to the backbone. This only results from poor design, or as part of a temporary work-around during changes to the network. A linkage occurs when two routers belong to the same area but are not physically connected to the same backbone area.

REVIEW QUESTIONS

1. Which of the following statements are true of Windows 2000 RRAS:
   a. RRAS supports nonpersistent connection through the use of demand-dial connections.
   b. Computers running RRAS can be configured to perform ICMP router discovery.
   c. RRAS is managed through Active Directory Users and Computers.
   d. RRAS supports VPN connections between routers.
   e. RRAS provides network address translation.
2. List the routing protocols included with Windows 2000 RRAS.

This scenario will be used in questions 3 through 7:

You are planning network modifications for a company with a central office in Dallas and branch offices in 50 locations in the United States. All offices are connected directly to the Internet via T-1 lines. Your group is designing the interoffice connectivity, which will involve using the Internet as the backbone for these connections. Several locations, including the central office, will have redundant connections to the ISP. It has been decided to use Windows 2000 RRAS servers dedicated to providing the routing services. You must provide for reliable, secure communications.

3. Which routing protocol will you include in your design?
4. Each branch office is using demand-dial connections. Someone has suggested using static routes on these routers rather than routing protocols. What would be gained by doing this? What is the biggest drawback to this strategy?
5. What two network connection devices will be needed for each T-1 connection?
6. One of the branch offices has both a T-1 and an ISDN connection. The RRAS router is connected to both interfaces as demand-dial interfaces. If you are using OSPF in your design, which connection route will be preferred? Explain your answer.
7. How will you limit your connection charges on the ISDN lines?
8. You are using Windows 2000 RRAS routers to route traffic on your private intranet. You have 12 subnets, with DHCP servers on two subnets and DHCP clients on 10 of the subnets. What is the recommended method for ensuring that all DHCP clients can receive the IP configuration from the DHCP server?
   a. Place the DHCP servers on RRAS routers.
   b. Configure the RRAS routers as DHCP relay agents.
   c. Place a DHCP server on each subnet.
   d. Configure a server on each network as a DHCP relay agent.
9. In a WAN connectivity design, you would like to secure the router-to-router communication with IPSec. What three rules must you keep in mind?

10. In a WAN connectivity design, you are considering using VPN tunnels on your routers for increased security. What two rules should you keep in mind?
11. If your design calls for restricting access to one of the locations to two groups, and placing time limits on one of the groups as well as different IP filters based on the group that is connecting, what will need to be configured?
12. Which would be a valid combination for providing authentication and encryption over a VPN?
    a. IPSec and PPTP
    b. ICMP and L2TP
    c. L2TP and IPSec
    d. RIP2 and OSPF
13. You need to set up a RRAS server as a remote access server. Some of the remote clients will be authenticating through smart cards. What authentication protocol must you implement on the RRAS server to allow for the smart card authentication of the remote clients?
14. Your company has 15 branch offices connecting to the central office through leased lines. Your CFO is concerned about the cost of these connections, especially since they are also paying for each office to have Internet connectivity. In general, what strategy would be more cost effective?
15. Your company is expanding rapidly and will soon outgrow its present office space. A study has indicated that 15% of the workforce could accomplish their work from home, saving the company the expense of maintaining office space. You have been assigned the task of designing connectivity for those workers who choose to telecommute. In addition to the computers that will be provided, what will each telecommuter need to work from home?
16. If a user's account properties are set to deny dial-in connections, but the remote access policy has a condition that allows connections from a group to which the user belongs, will the user be permitted to connect? Explain your answer.
17. Which of the following are remote access authentication protocols?
    a. EAP-TLS
    b. MPPE
    c. MS-CHAPv2
    d. PPP
    e. IPSec
18. Your IT manager has requested that the dial-in client authentication be centralized and that you set up an accounting system to track dial-in connections. What should you include in your remote access design to provide these services? Does such a solution come with Windows 2000?

7

19. What options do you have for enhancing a remote access design for availability and performance?
20. What three network protocols are supported for remote access clients?

Hands-on Projects

Project 7-1 Enabling RRAS Routing

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. You will use the Routing and Remote Access console to enable routing.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.
3. In the User Name box, type **administrator**.
4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)
6. Press **Return**, and when the desktop appears, click the **Start** button on the taskbar.
7. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**.
8. In the Tree pane, right-click your ***servername***.
9. Select **Configure and Enable Routing and Remote Access**. The Routing and Remote Access Server Setup wizard will appear.
10. Click the **Next** button. The Common Configurations page will appear.
11. Select **Network router**, and click the **Next** button.
12. In the Routed Protocols page, verify that TCP/IP is in the Protocols box, and then click the **Next** button. (If other protocols are listed, remove them after completing the wizard.)
13. In the Demand-Dial Connections page, select **Yes**, and then click the **Next** button.
14. In the IP Address Assignment page, select **Automatically**, and then click the **Next** button.
15. In the Completing the Routing and Remote Access Server Setup wizard, read the list of tasks that must be completed before using the router and list them below:

16. Click the **Finish** button. A Completing Installation message box will appear. In the Routing and Remote Access console, your server will have a green arrow in a white circle on its icon.
17. What objects appear in the Tree pane under your server? List them below:
18. Right-click your ***servername***. Notice that the option to Configure and Enable Routing and Remote Access is grayed out and Disable Routing and Remote Access is available.
19. Select **Properties**. Notice that router is selected, and the sub-setting LAN and demand-dial routing is selected.
20. Browse through the other tabs in this dialog box. You will see default settings, like Security, that you did not select through the wizard. But you will also see the setting you chose, such as the IP settings, in which IP routing is enabled, IP-based remote access and demand-dial connections are allowed, and the server will assign addresses using DHCP. You may modify any of these settings through the Properties dialog box.
21. Close the **Properties** dialog box.
22. If you plan to continue to Hands-on Project 7-2, leave the Routing and Remote Access console open; otherwise, close the console and log off.
23. If additional protocols (other than TCP/IP) appeared in Step 12, remove them now by using the Properties dialog box for each connection in Network and Dial-up Connections.

You have enabled LAN and demand-dial routing, but still have a few more tasks to do before you can use the router.

Project 7-2 Configuring RRAS for Demand-Dial

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. In Hands-on Project 7-1, you wrote down the steps that need to be completed to configure the router. In this project, you will add a demand-dial interface. Before you start the project, you will need the following information from your instructor:

- The IP address of the instructor's server
- A user account name that the instructor has added as dial-in credentials on a demand-dial interface on the instructor's server

- The domain in which this account is valid (all the computers in the class lab should be members of this domain)
- The password for this domain account

If the Routing and Remote Access console is still open on the desktop of your server, skip to Step 8; otherwise, start from the first step.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.
3. In the User Name box, type **administrator**.
4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)
6. Press **Return**, and when the desktop appears, click the **Start** button on the taskbar.
7. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**.
8. In the Tree pane, right-click **Routing Interfaces**.
9. Select **New Demand Dial Interface**. The Demand Dial Interface wizard will appear. Click **Next** on the Welcome page.
10. Give the interface a meaningful name; we suggest that you use the name of the router to which it connects. For this project use **Remote Router 1**. Click **Next**.
11. In the Connection Type page, select **Connect using VPN**. Then click the **Next** button.
12. In the VPN type page, you have three choices. Record these choices on the lines below:

    ____

    ____

    ____

13. Select **Automatic selection**, and then click the **Next** button.
14. In the Host name or IP address box, type the IP address of the instructor's computer. Click **Next**.
15. In the Protocols and Security page, ensure that **Route IP packets on this interface** is selected, and then click the **Next** button.
16. In the Dial Out Credentials box, enter the user name, domain, and password provided by your instructor. Then click the **Next** button.
17. Click the **Finish** button. The new interface will appear in the Routing Interfaces detail pane.
18. Right-click the **new interface** and select **Properties**.

19. On the General tab, verify that the interface has an IP address of the destination router (the instructor machine IP address).
20. Browse through the other four tab sheets to see all the options for properties.
21. What is the range of intervals that can be selected for demand dial idle time before hanging up? How would you allow an unlimited connection time?

22. What are you able to select in Dialing policy?

23. How would you modify the logon and encryption security settings?

24. Where would you modify the VPN and network settings?

25. Click the **Cancel** button in the Remote Router 1 Properties box.
26. If you will be continuing to the next project at this time, leave the Routing and Remote Access console open; otherwise, close all applications and log off.

In this project, you configured a demand-dial interface and explored the configuration options for a demand-dial interface.

Project 7-3 Adding Protocols to a Router

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. You will also need access to the Windows 2000 Advanced Server source files. Your instructor will make these available to you either on a share on the network or in a folder on your lab computer.

In this project, you will add the IPX/SPX protocols to the routing interface you added in the last project so that the routing support for these protocols can be seen in the Routing and Remote Access console. You would only add these protocols to the RRAS interfaces attached to networks on which these protocols were in use by clients and servers. You will also verify that your interfaces have addresses.

If the Routing and Remote Access console is still open on the desktop of your server, skip to Step 9; otherwise, start from the first step.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.
3. In the User Name box, type **administrator**.

4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar.
8. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**.
9. In the Details pane, right-click **Remote Router 1** and select **Properties**.
10. In the Properties dialog box, click the Networking tab, and then click the **Install** button.
11. Select **Protocol**, and then click the **Add** button.
12. In the list of network protocols, select **NWLink IPX/SPX**, and then click the **OK** button. At this point you may be prompted to provide the location of the Windows 2000 source files. If so, respond to the prompts until the properties dialog box for Remote Router 1 appears.

    In the components list, all installed components are listed. Those components with a check in the box are enabled. Those with a clear box are disabled for this interface.
13. Click the **Close** button.
14. Right-click the object for your server, and then click **Refresh**. In the **Routing and Remote Access** console, under the object for your server you will now see two objects for routing protocols. List them below:
15. Explore the nodes under IPX Routing where you will see RIP for IPX and SAP for IPX.
16. If you will be continuing to the next project at this time, leave the Routing and Remote Access console open; otherwise, close all applications and log off.

In this project, you added protocols to your network interface, and then viewed the related routing protocols in Routing and Remote Access. If you were setting up a router with several routing interfaces, you would now open the properties of each interface and remove the protocols that are not needed on that interface.

Project 7-4 Configuring IP Filters in RRAS

Hands-on Project

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. You will configure a router so that ICMP packets will not be passed in either direction on a single interface. If you are already logged on to your lab server with the Routing and Remote Access console open, you may skip to Step 9.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.

3. In the User Name box, type **administrator**.
4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar.
8. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**.
9. In the Tree pane, expand your ***servername***, expand **IP Routing**, and then click **General**.
10. In the Details pane, right-click the **Remote Router 1** interface you created in Project 7-2, and then select **Properties**.
11. On the General tab sheet, click the **Input Filters** button.
12. On the Input Filters page, click the **Add** button.
13. Click **Source** network to place a check in the box, and fill in the following information:
    - IP address: 192.168.1.0 (or an address supplied by your instructor)
    - Subnet mask: 255.255.255.0 (or a mask supplied by your instructor)
14. Click **Destination** network to place a check in the box, and fill in the following information:
    - IP address: 192.168.2.0 (or an address supplied by your instructor)
    - Subnet mask: 255.255.255.0 (or a mask supplied by your instructor)
15. Under Protocol, use the down-arrow button to select **ICMP**.
16. In the ICMP type box, type **8**.
17. In the ICMP code box, type **0** (zero), and then click the **OK** button.
18. In the Input Filters page, the new filter is listed, and the two radio buttons at the top are now active. The "Receive all packets except those that meet the criteria below" option is selected as the default. Leave this as the default. This input filter will now drop any ICMP Echo Request packets, blocking outside attempts to ping an internal interface.
19. Click the **OK** button to close the Input Filters page, and then click the **OK** button to close the Remote Router 1 Properties.
20. If you will be continuing to the next project at this time, leave the Routing and Remote Access console open; otherwise, close all applications and log off.

In this project, you configured a simple input filter on a routing interface to accept all traffic except the protocol, type, and code that you selected.

Project 7-5 Adding Static Routes to a Demand-Dial Interface for RRAS Routing

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. If you are already logged onto your lab server with the Routing and Remote Access console open, you may skip to Step 9.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.
3. In the User Name box, type **administrator**.
4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar.
8. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**.
9. In the Tree pane, expand ***servername***, expand **IP Routing**, right-click **Static Routes**, and then select **New Static Route**.
10. In the Static Route dialog box, select **Remote Router 1** for the interface.
11. In the Destination box, enter **192.168.20.0**.
12. In the Network mask box, enter **255.255.255.0**. Because this is a demand-dial interface, the gateway is grayed out because it is not configurable for a static route on a demand-dial interface.
13. In Metric, enter **4**.
14. Verify that the **Use this route to initiate demand-dial connections** check box is checked.
15. Click the **OK** button to close the Static Route dialog box.
16. If you will be continuing to the next project at this time, close the Routing and Remote Access console; otherwise, close all applications and log off.

You have configured a static route on the demand-dial interface. Now, any traffic for network 192.168.20.0 will be routed to the demand-dial interface.

Project 7-6 Creating a Remote Access Policy

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. If you are already logged on to your lab server you may skip to Step 8.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.

3. In the User Name box, type **administrator**.
4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar.
8. Select **Programs, Administrative Tools**, and click **Computer Management**.
9. In Computer Management, select **Local Users and Groups**.
10. Create the following group and users:
    - Group: "directsales"
    - User: Your first name. The password can be your choice.
11. In the Dial-in tab of the new user properties, select **Control access through Remote Access Policy**. Then click **OK**.
12. Add your new account to the direct sales group that you created.
13. Close Local Users and Groups, and then close the Computer Management console.
14. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**.
15. In the Tree pane, expand your ***servername***, right-click **Remote Access Policies**, and then select **New Remote Access Policy**.
16. In the Policy friendly name box of the Add Remote Access Policy page, type **DirectSales**, and then click the **Next** button.
17. On the Conditions page, click the **Add** button, select **Windows-Groups**, and then click the **Add** button.
18. In the Groups page, click the **Add** button, and then verify that the Look in box of the Select Groups page lists your *servername.*
19. Click **directsales**, click the **Add** button, and then click the **OK** button.
20. In the Groups page, click the **OK** button. The directsales group from your computer is now listed under Conditions. At this point you could click the **Add** button to select more conditions. We will not add additional conditions.
21. Click the **Next** button.
22. In the Permissions page, select what happens if a user attempting to make a connection matches the conditions.
23. Click the **Grant remote access permission** option button, and then click the **Next** button.
24. On the User Profile page, click the **Edit Profile** button.

25. On the Dial-in Constraints tab of the Edit Dial-in Profile dialog box, click the **Disconnect if idle for** check box, and enter **10** in the minutes box.
26. Click the **IP** tab, and then select **Server must supply an IP address**. Also notice that you can create IP packet filters with a remote access policy.
27. Click the **OK** button to complete your selection of profile settings.
28. In the User Profile page, click the **Finish** button.
29. In the Routing and Remote Access console, click **Remote Access Policies**. Notice that both the default policy, Allow access if dial-in permission is enabled, and your new policy, DirectSales, are listed. Now that you have a policy for the users who will be connecting to this RRAS server, you may delete the default policy or change the order in which they are processed for each connection.
30. Right-click the **DirectSales** policy, and then select **Move Up**. If the DirectSales policy conditions matches a connection attempt, no other policy will be evaluated for the connection.
31. If you will be continuing to the next project at this time, close the Routing and Remote Access console; otherwise, close all applications and log off.

In this project, you created a remote access policy that would restrict access to members of the DirectSales group. At the completion of the project, you moved the new remote access policy to the top of the list of remote access policies on your server, so that this policy will be evaluated first. If this is the only policy you wish to have applied to an RRAS server, you would delete the default policy.

Project 7-7 Implementing an IAS Server for a RADIUS Design

Hands-on Project

For this project, you will need a computer running Windows 2000 Server or Advanced Server connected to an IP network. You will also need access to the Windows 2000 Advanced Server source files. Your instructor will make these available to you either on a share on the network or in a folder on your lab computer.

In this project, you will install the Microsoft RADIUS server and configure that same server as a remote access server that is a client to the IAS RADIUS server. You will also configure IAS to log all authentication requests. If you are already logged on to your lab server, you may skip to Step 7.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Log On to Windows dialog box.
3. In the User Name box, type **administrator**.
4. In the Password box, type **password** (if this does not work, ask your instructor for the password).
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This will depend on the classroom configuration.)

6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar, and then select **Settings** and click **Control Panel**.
8. In Control Panel, double-click **Add/Remove Programs**, and then click the **Add/Remove Windows Components** button.
9. In the Components box of the Windows Components wizard, scroll down until **Networking Services** appears.

Caution

In the following steps, be very careful not to click the box. This would select all networking services, which is not a good thing!

10. Click the words **Networking Services**, and then click the **Details** button.
11. In the Networking Services page, check the **Internet Authentication Service** check box.
12. Click the **OK** button.
13. On the Windows Components page, click the **Next** button.
14. If prompted for files needed, enter the location for the server source files provided by your instructor, and then click the **OK** button.
15. In the Completing the Windows Components wizard, click the **Finish** button, and then close Add/Remove Programs and the Control Panel.
16. You will now find a new tool available from the Administrative Tools menu: Internet Authentication Service. Open this tool.
17. In the Tree pane, right-click **Internet Authentication Service (Local)**, and then click **Register Service in Active Directory**.
18. If a message appears that it is already registered in the Active Directory, click the **OK** button and continue with the next step.
19. In the Tree pane, right-click **Clients**, and then click **New Client**.
20. On the Name and Protocol page, type your ***servername*** in the Friendly name box, and then click the **Next** button.
21. On the Client Information page, type the address of your network adapter in the Client address (IP or DNS) box.
22. Also on the Client Information page in the Client-Vendor box, select **Microsoft**.
23. Still on the Client Information page, type **password** in the Shared secret and Confirm shared secret text boxes, and then click the **Finish** button.
24. Back in the Internet Authentication Service console, click **Remote Access Logging** in the Tree pane, right-click **Local File** in the details pane, and select **Properties**.

25. On the Settings tab, check the **Log authentication requests** and the **Log periodic status** check boxes, and then click the **Local File** tab.
26. Notice the settings available to you. The default location for the log file is in the log file directory. Look at that location and think of a reason why you would want to change that location. Write your explanation below:

    ____

    ____

27. Click the **OK** button on the Local File Properties page and close the **Internet Authentication Service** console.
28. Select **Programs**, **Administrative Tools**, and click **Routing and Remote Access**. The Routing and Remote Access console may already be open.
29. Right-click your ***servername*** and select **Properties**.
30. On the General tab of the Properties dialog box, check the **Remote access server** check box, and then click the **OK** button. A Routing and Remote Access message appears warning that the router must be restarted. Surprisingly, this is *just* the router that must restart. It will not cause your computer to restart, so click the **Yes** button and wait several minutes while the RRAS service is first started, and then restarted.
31. In the Routing and Remote Access console, right-click your ***servername***, and then click **Properties**.
32. In the Properties dialog box, click the **Security** tab.
33. In the Authentication provider box, select **RADIUS Authentication**, and then click the **Configure** button.
34. In the RADIUS Authentication page, click the **Add** button.
35. In the Add RADIUS Server box, type your ***servername*** in the Server name box, and then click the **Change** button next to Secret.
36. Type **password** in the New secret and Confirm new secret text boxes, and then click the **OK** button. This secret is used by the RADIUS server and the client (the remote access server) when they authenticate a secure channel for communications between the RADIUS server and RADIUS client.
37. Click **OK** in the Add RADIUS Server and RADIUS Authentication boxes.
38. A box will appear, warning that the Routing and Remote Access service must be restarted before the RADIUS setting will take effect. Click the **OK** button in the warning box.
39. Click **OK** in the server Properties box, read the warning that appears, and click the **Yes** button in the warning box.

40. In the Routing and Remote Access console, notice that the remote access policies have been removed. When a Windows 2000 remote access server becomes a RADIUS client to IAS, the remote access policies are moved to the IAS server.
41. Open the IAS console and confirm that the remote access policies were moved here, including the policy you created in Hands-on Project 7-6.

In this project, you installed and configured Internet Authentication Service (IAS) and then made the RRAS remote access server on your server a RADIUS client to IAS. In your network design, you are more likely to place these services on separate servers and have many RRAS remote servers configured as RADIUS clients to the IAS server.

CASE PROJECTS

Case 7-1 Creating a Router Design

ASDFG is a manufacturing company based in San Jose with manufacturing and warehouse centers in Vancouver, Montreal, and Detroit. They presently connect the branch offices directly to San Jose with leased lines. In addition, the main office has Internet connections, but only two of the other sites do.

Two of their suppliers have moved their business-to-business product ordering system to the Internet, and both are offering large incentives for the first year of usage. This will require reliable Internet connections to ISPs from all locations. After load testing, it was decided that San Jose will have a T-3 connection and the three branch offices will have T-1 connections. Once these are in place, the network manager would like to move their interoffice connections to an Internet-based model. The San Jose location has four subnets, including one where the users are accessing extremely confidential data that should not be accessible from any other part of the network, although these same users must be able to access servers in other parts of the network.

Create a router design for this scenario, providing interconnectivity between all the offices. It has been decided that, in addition to the T-carrier connections at each site, there should be redundant connections to the Internet at lesser speed and bandwidth.

Do the following:

1. Draw a diagram showing where you would place routers and the connection devices.
2. Write a description that includes placement of servers, transport protocol(s), routing protocols, enhancements for security, enhancements for availability and performance, and options for isolating the subnet that has confidential data.

Case 7-2 Designing a Remote Access Solution

ASDFG has a mobile sales force of 50 people who previously dialed into remote access servers in the central office from their notebook computers running Windows 98. This involved high connection charges. You are involved in designing a remote access solution in which they will dial in over the Internet. They need e-mail and access to a central inventory and order-entry system in San Jose. Their connections must be secure.

Write a description outlining the components needed for these connections, accompanied by a drawing of a single dial-in connection and components. Be sure to include the placement of servers and the transport protocol(s). In addition, you should provide enhancements for security, availability, and performance.

Case 7-3 Outlining a RADIUS Solution

ASDFG has purchased another company. This adds several additional remote manufacturing and warehousing locations. They must give remote access to the San Jose private network to employees at these new locations, but they are very concerned about security. They also need to establish accounting of remote access usage.

Write a description outlining how you would use a RADIUS solution in your design. Draw a graphic to represent your solution. Include the placement of participating servers and other added components.